

Subscribe (Full Service) Register (Limited Service, Free) Login

Search: The ACM Digital Library O The Guide

query optimization bitmap

SEARCH

THE ACTIONOMICAL LIBRARY

Feedback Report a problem Satisfaction survey

Terms used <u>query optimization</u> <u>bitmap</u>

Found 15,520 of 198,617

Sort results by

publication date

Save results to a Binder ? Search Tips

Try an Advanced Search Try this search in The ACM Guide

Display results

expanded form

Open results in a new window

Results 101 - 120 of 200 Best 200 shown

Result page: previous 1 2 3 4 5 6 7 8 9 10

Relevance scale

101 Streaming XML: XPath queries on streaming data

Feng Peng, Sudarshan S. Chawathe

June 2003 Proceedings of the 2003 ACM SIGMOD international conference on Management of data SIGMOD '03

Publisher: ACM Press

Full text available: pdf(433.73 KB)

Additional Information: full citation, abstract, references, citings, index terms

We present the design and implementation of the XSQ system for querying streaming XML data using XPath 1.0. Using a clean design based on a hierarchical arrangement of pushdown transducers augmented with buffers, XSQ supports features such as multiple predicates, closures, and aggregation. XSQ not only provides high throughput, but is also memory efficient: It buffers only data that must be buffered by any streaming XPath processor. We also present an empirical study of the performance character ...

102 Online analytic processing (OLAP): Spreadsheets in RDBMS for OLAP



Andrew Witkowski, Srikanth Bellamkonda, Tolga Bozkaya, Gregory Dorman, Nathan Folkert, Abhinav Gupta, Lei Shen, Sankar Subramanian

June 2003 Proceedings of the 2003 ACM SIGMOD international conference on Management of data SIGMOD '03

Publisher: ACM Press

Full text available: pdf(182.22 KB)

Additional Information: full citation, abstract, references, citings, index terms

One of the critical deficiencies of SQL is lack of support for n-dimensional array-based computations which are frequent in OLAP environments. Relational OLAP (ROLAP) applications have to emulate them using joins, recently introduced SQL Window Functions [18] and complex and inefficient CASE expressions. The designated place in SQL for specifying calculations is the SELECT clause, which is extremely limiting and forces the user to generate queries using nested views, subqueries and complex joins ...

103 XML indexing and compression: Efficient processing of joins on set-valued attributes



Nikos Mamoulis

June 2003 Proceedings of the 2003 ACM SIGMOD international conference on Management of data SIGMOD '03

Publisher: ACM Press

Full text available: pdf(678.13 KB)

Additional Information: full citation, abstract, references, citings, index terms

Object-oriented and object-relational DBMS support set valued attributes, which are a natural and concise way to model complex information. However, there has been limited research to-date on the evaluation of query operators that apply on sets. In this paper we study the join of two relations on their set-valued attributes. Various join types are considered, namely the set containment, set equality, and set overlap joins. We show that the inverted file, a powerful index for selection queries, c ...

104 Statistics: Extended wavelets for multiple measures

Antonios Deligiannakis, Nick Roussopoulos

June 2003 Proceedings of the 2003 ACM SIGMOD international conference on Management of data SIGMOD '03

Publisher: ACM Press

Full text available: pdf(322.18 KB)

Additional Information: full citation, abstract, references, citings, index terms

While work in recent years has demonstrated that wavelets can be efficiently used to compress large quantities of data and provide fast and fairly accurate answers to queries, little emphasis has been placed on using wavelets in approximating datasets containing multiple measures. Existing decomposition approaches will either operate on each measure individually, or treat all measures as a vector of values and process them simultaneously. We show in this paper that the resulting *individual*

105 Industrial track session 2: server technology: Multi-dimensional clustering: a new

data layout scheme in DB2

Sriram Padmanabhan, Bishwaranjan Bhattacharjee, Tim Malkemus, Leslie Cranston, Matthew Huras

June 2003 Proceedings of the 2003 ACM SIGMOD international conference on Management of data SIGMOD '03

Publisher: ACM Press

Full text available: pdf(168.35 KB) Additional Information: full citation, abstract, references, index terms

We describe the design and implementation of a new data layout scheme, called multi-dimensional clustering, in DB2 Universal Database Version 8. Many applications, e.g., OLAP and data warehousing, process a table or tables in a database using a multi-dimensional access paradigm. Currently, most database systems can only support organization of a table using a primary clustering index. Secondary indexes are created to access the tables when the primary key index is not applicable. Unfortunately, ...

106 Discovering all most specific sentences

Dimitrios Gunopulos, Roni Khardon, Heikki Mannila, Sanjeev Saluja, Hannu Toivonen, Ram Sewak Sharma

June 2003 ACM Transactions on Database Systems (TODS), Volume 28 Issue 2

Publisher: ACM Press

Full text available: pdf(283.09 KB)

Additional Information: full citation, abstract, references, citings, index terms

Data mining can be viewed, in many instances, as the task of computing a representation of a theory of a model or a database, in particular by finding a set of maximally specific sentences satisfying some property. We prove some hardness results that rule out simple approaches to solving the problem. The *a priori* algorithm is an algorithm that has been successfully applied to many instances of the problem. We analyze this algorithm, and prove that is optimal when the maximally specific sen ...

Keywords: Data mining, association rules, learning with membership queries, maximal frequent sets, minimal keys

107 Information retrieval 2: Dynamic maintenance of web indexes using landmarks



Lipyeow Lim, Min Wang, Sriram Padmanabhan, Jeffrey Scott Vitter, Ramesh Agarwal May 2003 Proceedings of the 12th international conference on World Wide Web **WWW '03**



Publisher: ACM Press

Full text available: pdf(233.78 KB)

Additional Information: full citation, abstract, references, citings, index terms

Recent work on incremental crawling has enabled the indexed document collection of a search engine to be more synchronized with the changing World Wide Web. However, this synchronized collection is not immediately searchable, because the keyword index is rebuilt from scratch less frequently than the collection can be refreshed. An inverted index is usually used to index documents crawled from the web. Complete index rebuild at high frequency is expensive. Previous work on incremental inverted in ...

Keywords: inverted files, update processing

108 Astrolabe: A robust and scalable technology for distributed system monitoring,





management, and data mining

Robbert Van Renesse, Kenneth P. Birman, Werner Vogels

May 2003 ACM Transactions on Computer Systems (TOCS), Volume 21 Issue 2

Publisher: ACM Press

Full text available: pdf(341.62 KB)

Additional Information: full citation, abstract, references, citings, index terms

Scalable management and self-organizational capabilities are emerging as central requirements for a generation of large-scale, highly dynamic, distributed applications. We have developed an entirely new distributed information management system called Astrolabe. Astrolabe collects large-scale system state, permitting rapid updates and providing on-the-fly attribute aggregation. This latter capability permits an application to locate a resource, and also offers a scalable way to track sys ...

Keywords: Aggregation, epidemic protocols, failure detection, gossip, membership, publish-subscribe, scalability

109 Data page layouts for relational databases on deep memory hierarchies

Anastassia Ailamaki, David J. DeWitt, Mark D. Hill

November 2002 The VLDB Journal — The International Journal on Very Large Data Bases, Volume 11 Issue 3

Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(593.86 KB) Additional Information: full citation, abstract, index terms

Relational database systems have traditionally optimized for I/O performance and organized records sequentially on disk pages using the N-ary Storage Model (NSM) (a.k.a., slotted pages). Recent research, however, indicates that cache utilization and performance is becoming increasingly important on modern platforms. In this paper, we first demonstrate that in-page data placement is the key to high cache performance and that NSM exhibits low cache utilization on modern platforms. Next, we ...

Keywords: Cache-conscious database systems, Disk page layout, Relational data placement

110 Strategies for processing ad hoc queries on large data warehouses

Kurt Stockinger, Kesheng Wu, Arie Shoshani

November 2002 Proceedings of the 5th ACM international workshop on Data

Warehousing and OLAP DOLAP '02

Publisher: ACM Press

Full text available: pdf(245.31 KB)

Additional Information: full citation, abstract, references, citings, index terms

As data warehousing applications grow in size, existing data organizations and access strategies, such as relational tables and B-tree indexes, are becoming increasingly ineffective. The two primary reasons for this are that these datasets involve many attributes and the queries on the data usually involve conditions on small subsets of the attributes. Two strategies are known to address these difficulties well, namely vertical partitioning and bitmap indexes. In this paper, we summarize our exp ...

111 Web services and performance evaluation: Indexing web access-logs for pattern



queries

Alexandros Nanopoulos, Yannis Manolopoulos, Maciej Zakrzewicz, Tadeusz Morzy November 2002 Proceedings of the 4th international workshop on Web information and data management WIDM '02

Publisher: ACM Press

Full text available: pdf(187.24 KB)

Additional Information: full citation, abstract, references, citings, index terms

In this paper, we develop a new indexing method for large web access-logs. We are concerned with pattern queries, which advocate the search for access sequences that contain certain query patterns. This kind of queries find applications in processing web-log mining results (e.g., finding typical/atypical access-sequences). The proposed method focuses on scalability to web-logs' sizes. For this reason, we examine the gains due to signature-trees, which can further improve the scalability to very ...

112 Adding subqueries to MySQL, what does it take to have a decision-support engine?



Antonio Badia, Matt Chanda, Bin Cao

November 2002 Proceedings of the 5th ACM international workshop on Data Warehousing and OLAP DOLAP '02

Publisher: ACM Press

Full text available: pdf(200.24 KB) Additional Information: full citation, abstract, references, index terms

MySQL is an OLTP system without the ability to handle SQL subqueries. As part of our project in query optimization, we add subquery processing to MySQL through rewriting. First, we incorporate the ability to handle subqueries in the FROM clause; then rewrite queries with subqueries in the WHERE clause by moving the subquery to the FROM clause and introducing suitable conditions to link the subquery and the main query tables. The result is a complete and correct unnesting procedure for SQL querie ...

Keywords: SLQ, query optimization, unnesting

113 Efficient similarity search for market basket data

Alexandros Nanopoulos, Yannis Manolopoulos

October 2002 The VLDB Journal — The International Journal on Very Large Data Bases, Volume 11 Issue 2

Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(212.88 KB) Additional Information: full citation, abstract, citings, index terms

Several organizations have developed very large market basket databases for the maintenance of customer transactions. New applications, e.g., Web recommendation systems, present the requirement for processing similarity queries in market basket databases. In this paper, we propose a novel scheme for similarity search queries in basket data. We develop a new representation method, which, in contrast to existing approaches, is proven to provide correct results. New algorithms are proposed for the ...

Keywords: Data mining, Market basket data, Nearest-neighbor, Similarity search

114 Exploiting common subqueries for complex query optimization

Yingying Tao, Qiang Zhu, Calisto Zuzarte

September 2002 Proceedings of the 2002 conference of the Centre for Advanced Studies on Collaborative research CASCON '02

Publisher: IBM Press

Full text available: pdf(300.77 KB) Additional Information: full citation, abstract, references, index terms

As database technology is applied to more and more application areas, user queries on a database become more and more complex. Existing query optimization techniques were not developed for dealing with complex queries and may suffer from some serious problem such as intolerably long optimization time and poor optimizing results. To tackle this challenge, we introduce a new technique to improve the quality of complex query optimization in this paper. The key idea is to exploit the common subqueri ...

115 Frequent patterns I: DualMiner: a dual-pruning algorithm for itemsets with constraints



Cristian Bucila, Johannes Gehrke, Daniel Kifer, Walker White July 2002 Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining KDD '02

Publisher: ACM Press

Full text available: pdf(1.22 MB)

Additional Information: full citation, abstract, references, citings, index terms

Constraint-based mining of itemsets for questions such as "find all frequent itemsets where the total price is at least \$50" has received much attention recently. Two classes of constraints, monotone and antimonotone, have been identified as very useful. There are algorithms that efficiently take advantage of either one of these two classes, but no previous algorithms can efficiently handle both types of constraints simultaneously. In this paper, we present the first algorithm (called Dua ...

116 Research sessions: implementation techniques: Fractal prefetching B±-Trees:



optimizing both cache and disk performance

Shimin Chen, Phillip B. Gibbons, Todd C. Mowry, Gary Valentin

June 2002 Proceedings of the 2002 ACM SIGMOD international conference on Management of data SIGMOD '02

Publisher: ACM Press

Full text available: pdf(1.49 MB)

Additional Information: full citation, abstract, references, citings, index terms

B+-Trees have been traditionally optimized for I/O performance with disk pages as tree nodes. Recently, researchers have proposed new types of B+-Trees optimized for CPU cache performance in main memory environments, where the tree node sizes are one or a few cache lines. Unfortunately, due primarily to this large discrepancy in optimal node sizes, existing disk-optimized B+-Trees suffer from poor cache performance while cacheoptimized B+-Trees exhibi ...

117 Research sessions: continuous queries and streams: Continuously adaptive



continuous queries over streams

Samuel Madden, Mehul Shah, Joseph M. Hellerstein, Vijayshankar Raman June 2002 Proceedings of the 2002 ACM SIGMOD international conference on Management of data SIGMOD '02

Publisher: ACM Press

Full text available: pdf(1.59 MB)

Additional Information: full citation, abstract, references, citings, index terms

We present a continuously adaptive, continuous query (CACQ) implementation based on the eddy query processing framework. We show that our design provides significant performance benefits over existing approaches to evaluating continuous queries, not only because of its adaptivity, but also because of the aggressive cross-query sharing of work and space that it enables. By breaking the abstraction of shared relational algebra expressions, our Telegraph CACQ implementation is able to share physica ...

118 Research sessions: query processing II: Efficient k-NN search on vertically

decomposed data

Arjen P. de Vries, Nikos Mamoulis, Niels Nes, Martin Kersten

June 2002 Proceedings of the 2002 ACM SIGMOD international conference on Management of data SIGMOD '02

Publisher: ACM Press

Full text available: pdf(1.26 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> terms

Applications like multimedia retrieval require efficient support for similarity search on large data collections. Yet, nearest neighbor search is a difficult problem in high dimensional spaces, rendering efficient applications hard to realize: index structures degrade rapidly with increasing dimensionality, while sequential search is not an attractive solution for repositories with millions of objects. This paper approaches the problem from a different angle. A solution is sought in an unconvent ...

119 Research sessions: query processing I: Partial results for online query processing

Vijayshankar Raman, Joseph M. Hellerstein

June 2002 Proceedings of the 2002 ACM SIGMOD international conference on Management of data SIGMOD '02

Publisher: ACM Press

Full text available: pdf(1.45 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> terms

Traditional query processors generate full, accurate query results, either in batch or in pipelined fashion. We argue that this strict model is too rigid for exploratory queries over diverse and distributed data sources, such as sources on the Internet. Instead, we propose a looser model of querying in which a user submits a broad initial query outline, and the system continually generates *partial* result tuples that may contain values for only some of the output fields. The user can watch ...

120 Research sessions: implementation techniques: Implementing database operations





using SIMD instructions

Jingren Zhou, Kenneth A. Ross

June 2002 Proceedings of the 2002 ACM SIGMOD international conference on Management of data SIGMOD '02

Publisher: ACM Press

Full text available: pdf(1.39 MB)

Additional Information: <u>full citation</u>, <u>abstract</u>, <u>references</u>, <u>citings</u>, <u>index</u> terms

Modern CPUs have instructions that allow basic operations to be performed on several data elements in parallel. These instructions are called SIMD instructions, since they apply a single instruction to multiple data elements. SIMD technology was initially built into commodity processors in order to accelerate the performance of multimedia applications. SIMD instructions provide new opportunities for database engine design and implementation. We study various kinds of operations in a database con ...

Results 101 - 120 of 200 Result page: previous 1 2 3 4 5 6 7 8 9 10 next

The ACM Portal is published by the Association for Computing Machinery. Copyright @ 2007 ACM, Inc.

Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat Q QuickTime Windows Media Player Real Player





March 16, 2007

USPTO

Search

Full Text

Concept

Document ID

Recent Disclosures

Displaying records #11 through 20 out of 364

Result # 11 Relevance: OOO

Process for Representing Screen-Sized Graphics as an Icon

1994-03-01 IPCOM000111726D

English

The invention is a process for selectively converting either a vector representation of a graphic or, alternately, a bit-map representation of a screen-sized graphic into an icon.

Other

Prior Art Home

Support

Logout

About IP.com

FAQs

Contact Us

Relevance: OOO 3 6

of Using Compressed Bit-Map for Memory Utilization

IPCOM000120439D

English

ie invented compressed bit map as a better way of keeping track of the b

Result # 13 Relevance: 〇〇〇 にかご

Improved Graphic Image Bitmap Structure for Personal Computers

1994-07-01

IPCOM000113081D

English

Described is an architectural implementation designed to improve the graphic image bit and controlling apparatus and to enable horizontal, vertical, and quadrangle area acces Personal Computers (PCs). The implementation relates to the graphic ...

Result # 14 Relevance: OOO

Method to Dynamically Allocate Bitmaps on a Per Page Basis for Perfo Improvement in the Adobe PDF Rip Process

2006-03-08

IPCOM000134525D

English

Disclosed is a method(s) to force the Adobe* Configurable PostScript** Interpreter (i.e render into a new raster bitmap for each output bitmap it produces when interpreting e (Portable Document Format) or PostScript language programs. The reason for ...

Result # 15 Relevance: 〇〇〇〇〇〇〇〇

Inter-Program Bitmap Data Distribution

1989-11-01

IPCOM000036909D

English

This article describes a method for interchanging bitmaps in a mul tasking environment independent applications which requires minimal system overhead and minimizes progr to interchange bitmap data between applications. Inter-Program Bitmap Data ...

Result # 16 Relevance: OOO

Extracting Data From Bitmap in Display

1990-02-01

IPCOM000099662D

English

In graphics displays, data to be operated on by a graphics engine may be stored in a pibitmap. This disclosure is concerned with extracting data from a single plane of a pixel another pixel bitmap. The technique described enables only selected bits ...

Result # 17 Relevance: QQ 1. 60.

Recursive Dilation for Visual ICON Selection

1994-04-01

IPCOM000111985D

English

A method is presented for an efficient process for providing a technique for CUA's defin

Visual for Icon Selection.

Result # 18 Relevance: QQ

Hand-Held CRT Scanner

1990-06-01 IPCOM000100907D English

A solid-state camera and a motor-driven X-Y table are used with the distortion measure shown in Fig. 1. Test patterns displayed on the CRT are digitized with this equipment se map of the image is created for data manipulation. Seven minutes is required ...

Result # 19 Relevance: 🔾 🔾

Time-Lapse Controlled Bitmap Display of Complex Graphical Data

1994-06-01

IPCOM000112652D

Described is a method of displaying a complex graphic image on a Personal Computer (windows-based application. A bitmap image is built in store and called to display it as r redraws avoiding repeated creations from scratch. While an application is ...

Relevance: QQ Result # 20

AUTOMATED REGRESSION TESTING USING A VIRTUAL PRINTER

IPCOM000013021D 2000-08-01

Disclosed is a method for fully automating the traditionally manual process of regressio printer management software. The following scenarios are addressed, which previously manual intervention: hardware or physical media failures; simultaneous disablement ...

Displaying page 2 of 37 << FIRST | < BACK | NEXT > | LAST >>

Search query: bitmap

New search | Modify this search | Search within current results

Copyright @ 2007 IP.com, Inc. All rights reserved. |





Copyright © 2007 IP.com, Inc. All rights reserved. |

March 16, 2007

USPTO

Full Text	Displaying records #1 through 1 out of 1
Concept	Result # 1 Relevance: CCCCC
Document ID	DISTRIBUTED QUERY PROCESSING IN LOCAL - NETWORK DATABASES
Recent Disclosures	1984-12-31 IPCOM000128335D English to Distributed Databases
Other Prior Art Home	Purpose of the Thesis . , 1.4 Thesis Plan page 2 Distributed
	Displaying page 1 of 1 << FIRST < BACK NEXT > LAST >>
Support	pispidying page 1 of 1 (Clind) Conce Next > Dad >>
	- Displaying page 1 of 1 CCTRST CACK NEXT > DAST >>
Logout	Search query: query optimization and bitmap